

CLAIMS:

What is claimed is:

5 1. A method of transmitting data packets from a system area network device to an external network device, comprising:
passing data generated by a host process to a host channel adapter; and

10 passing the data from the host channel adapter directly to a router coupled to an external network.

15 2. The method of claim 1, wherein passing the data generated by a host process to a host channel adapter includes invoking an Internet Protocol (IP) over InfiniBand (IB) device driver.

20 3. The method of claim 2, wherein passing data generated by a host process to a host channel adapter includes creating an IP over IB Queue Pair in the host channel adapter for use with the IP over IB device driver.

25 4. The method of claim 2, wherein the step of passing data generated by a host process to a host channel adapter is performed in response to an I/O Transmit transaction being received by the IP over IB device driver.

30 5. The method of claim 4, wherein the I/O Transmit transaction originates from one of a user level program and a kernel level program.

PCT/US2007/038812

6. The method of claim 4, wherein the I/O Transmit transaction includes one or more pointers to one or more memory regions which contain the data, and wherein the

5 I/O Transmit transaction further includes one of a destination address and destination address handle.

7. The method of claim 1, wherein passing data generated by a host process to a host channel adapter

10 includes using a Post Send verb to instruct the host channel adapter to send data from system memory to a designated destination.

8. The method of claim 1, wherein the data is passed to the host channel adapter as one of a Raw Datagram and a Unreliable Datagram.

15

9. An apparatus for transmitting data packets from a system area network device to an external network device,

20 comprising:

means for passing data generated by a host process to a host channel adapter; and

means for passing the data from the host channel adapter directly to a router coupled to an external

25 network.

10. The apparatus of claim 9, wherein the means for passing the data generated by a host process to a host channel adapter includes means for invoking an Internet

30 Protocol (IP) over InfiniBand (IB) device driver.

PCT/US2009/038005

11. The apparatus of claim 10, wherein the means for
passing data generated by a host process to a host
channel adapter includes means for creating an IP over IB
Queue Pair in the host channel adapter for use with the
5 IP over IB device driver.

12. The apparatus of claim 10, wherein the means for
passing data generated by a host process to a host
channel adapter operates in response to an I/O Transmit
10 transaction being received by the IP over IB device
driver.

13. The apparatus of claim 12, wherein the I/O Transmit
transaction originates from one of a user level program
15 and a kernel level program.

14. The apparatus of claim 12, wherein the I/O Transmit
transaction includes one or more pointers to one or more
memory regions which contain the data, and wherein the
20 I/O Transmit transaction further includes one of a
destination address and destination address handle.

15. The apparatus of claim 9, wherein the means for
passing data generated by a host process to a host
25 channel adapter includes means for using a Post Send verb
to instruct the host channel adapter to send data from
system memory to a designated destination.

16. The apparatus of claim 9, wherein the data is passed
30 to the host channel adapter as one of a Raw Datagram and
a Unreliable Datagram.

17. A computer program product in a computer readable medium for transmitting data packets from a system area network device to an external network device, comprising:

first instructions for passing data generated by a host process to a host channel adapter; and

second instructions for passing the data from the host channel adapter directly to a router coupled to an external network.

10 18. The computer program product of claim 17, wherein
the first instructions for passing the data generated by
a host process to a host channel adapter include
instructions for invoking an Internet Protocol (IP) over
InfiniBand (IB) device driver.

15 19. The computer program product of claim 18, wherein
the first instructions for passing data generated by a
host process to a host channel adapter include
instructions for creating an IP over IB Queue Pair in the
20 host channel adapter for use with the IP over IB device
driver.

20. The computer program product of claim 18, wherein
the first instructions for passing data generated by a
host process to a host channel adapter are executed in
response to an I/O Transmit transaction being received by
the IP over IB device driver.

21. The computer program product of claim 20, wherein
30 the I/O Transmit transaction originates from one of a
user level program and a kernel level program.

卷之三

22. The computer program product of claim 20, wherein the I/O Transmit transaction includes one or more pointers to one or more memory regions which contain the data, and wherein the I/O Transmit transaction further 5 includes one of a destination address and destination address handle.

23. The computer program product of claim 17, wherein the first instructions for passing data generated by a 10 host process to a host channel adapter include instructions for using a Post Send verb to instruct the host channel adapter to send data from system memory to a designated destination.

15 24. The computer program product of claim 17, wherein the data is passed to the host channel adapter as one of a Raw Datagram and a Unreliable Datagram.

20 25. A method of routing data between a system area network and an external network, comprising:
receiving data;
parsing a routing header of the data;
identifying an output port of the router based on the parsing of the routing header; and
25 sending the data out of the router via the identified output port.

26. The method of claim 25, wherein identifying an output port of the router includes examining one of an 30 InfiniBand Global Router Header's Destination Global Identifier and an IPv6 Destination Address.

27. The method of claim 25, wherein if the data is an Unreliable Datagram and the identified output port is not an InfiniBand output port, an InfiniBand Transport Header associated with the data is discarded.

5

28. The method of claim 25, wherein sending the data out of the router includes creating an InfiniBand link layer header for the data.

10 29. The method of claim 28, wherein the InfiniBand link layer header identifies a host channel adapter receive queue.

15 30. The method of claim 28, wherein the InfiniBand link layer header identifies an external network.

31. A computer program product in a computer readable medium for routing data between a system area network and an external network, comprising:

20 first instructions for receiving data;
second instructions for parsing a routing header of the data;

third instructions for identifying an output port of the router based on the parsing of the routing header;

25 and

fourth instructions for sending the data out of the router via the identified output port.

32. The computer program product of claim 31, wherein
30 the third instructions for identifying an output port of the router include instructions for examining one of an InfiniBand Global Router Header's Destination Global

Identifier and an IPv6 Destination Address.

33. The computer program product of claim 31, wherein if the data is an Unreliable Datagram and the identified 5 output port is not an InfiniBand output port, an InfiniBand Transport Header associated with the data is discarded.

34. The computer program product of claim 31, wherein 10 the fourth instructions for sending the data out of the router include instructions for creating an InfiniBand link layer header for the data.

35. The method of claim 34, wherein the InfiniBand link 15 layer header identifies a host channel adapter receive queue.

36. The method of claim 34, wherein the InfiniBand link layer header identifies an external network. 20

37. An apparatus for routing data between a system area network and an external network, comprising:
means for receiving data;
means for parsing a routing header of the data;
25 means for identifying an output port of the router based on the parsing of the routing header; and
means for sending the data out of the router via the identified output port.

30 38. The apparatus of claim 37, wherein the means for identifying an output port of the router includes means for examining one of an InfiniBand Global Router Header's

Destination Global Identifier and an IPv6 Destination Address.

39. The apparatus of claim 37, wherein if the data is an Unreliable Datagram and the identified output port is not an InfiniBand output port, an InfiniBand Transport Header associated with the data is discarded.
40. The apparatus of claim 37, wherein the means for sending the data out of the router includes creating an InfiniBand link layer header for the data.
41. The apparatus of claim 40, wherein the InfiniBand link layer header identifies a host channel adapter receive queue.
42. The apparatus of claim 40, wherein the InfiniBand link layer header identifies an external network.